

Compression with Actions

Lei Zhao, Yeow-Khiang Chia and Tsachy Weissman

Abstract---We consider the setting where actions can be used to modify a state sequence before compression. The minimum rate needed to losslessly describe the optimal modified sequence is characterized when the state sequence is either non-causally or causally available at the action encoder. The achievability is closely related to the optimal channel coding strategy for channel with states. We also extend the analysis to the the lossy case.

I. INTRODUCTION

Consider the standard Shannon-theoretic lossy source coding setting where we have a source S^n that we wish to perform lossy compression on. The encoder receives the source S^n and produces an index M that is sent to the decoder. Based on the index, the decoder produces a lossy reconstruction, \hat{S}^n , such that the per symbol distortion constraint is satisfied. An alternative view of this problem is as one where the encoder is first required to produce the reconstruction sequence \hat{S}^n and then uses a lossless compression algorithm to describe \hat{S}^n to the decoder. This point of view on lossy source coding, depicted in Figure 1, has been instrumental in recent developments of an approach to universal and implementable lossy compressors, cf. [14], [15] and references therein.

In this paper, we generalize the above setting by asking the following question: what if the encoder makes an “error” in outputting the reconstruction sequence? The encoder may wish to take “actions” to output a sequence of reconstruction symbols. However, due to noise, the reconstruction symbols at the output of the encoder may be different from the intended reconstruction symbols. In this case, we are still interested in sending the reconstruction sequence (the *modified* source sequence) to the decoder. The question then is, what is the optimal rate-distortion tradeoff in such a scenario? As a more concrete example, consider lossy compression of a binary source S^n . With the source as input, the encoder first attempts to output the desired reconstruction sequence, but due to errors in the circuitry of the encoder, a bit that is meant to be one can still be zero with some probability and vice versa. Using a universal lossless compression algorithm, we transmit the output of the “faulty” encoder, \hat{S}^n , to the decoder. We are now interested in the optimum rate-distortion tradeoff under the assumption of a “faulty” encoder.

As another example of our general setting, which may seem at first sight to be unrelated to the question we asked above, imagine that we have a number of robots working on a factory




Fig. 1. Lossy source coding.

floor and the positions of all the robots need to be reported to a remote location. Letting S represent the position of a robot, we would expect to send $H(S)$ bits to the remote location. However, what if the robots can take *actions* to change their positions so that they can be more efficiently described? A local command center can first give commands (actions) to the robots so that they move in a cooperative way into a final position sequence that requires fewer bits to describe. The command center may face two issues in general: cost constraints and uncertainty. A cost constraint occurs because each robot should save its power and not move too far away from its current location. The uncertainty is a result of the robots not moving exactly as instructed by the local command center.

Both examples are instantiations of the problem setting illustrated in Fig. 2 (Formal definitions are given in the next section). Here, S^n is our observed source (or state) sequence. We assume a general cost function $\Lambda(a, s, y)$ and a general relation, specified by a conditional PMF $p(y|a, s)$, relating the modified source sequence to be compressed to the original source sequence (state) and action taken by the encoder toward modifying it. As shown in the preceding examples, we are interested in compressing the final output Y^n .




Fig. 2. Compression with actions. The Action encoder first observes the state sequence S^n and then generates an action sequence A^n . The i th output Y_i is the output of a channel $p(y|a, s)$ when $a = A_i$ and $s = S_i$. The compressor outputs a description of Y^n , $M \in [1 : 2^{nR}]$, from Y^n alone if the side information Z^n is not available at the compressor. If the side information is available, then the compressor generates the description based on Y^n and Z^n . The remote decoder generates \hat{Y}^n based on M and its available side information Z^n as a reconstruction of Y^n .

Our problem setup is also closely related to the channel coding problem when the state information is available at the encoder. The case where the state information is causally available was first solved by Shannon in [4]. When the state information is non-causally known at the encoder, the channel capacity result was derived in [2] and [3]. Various interesting extensions can be found in [5]–[9]. The difference in our approach described here is that we make the output of the channel as compressible as possible. Our main results when the decoder requires lossless reconstruction are given in sec-

Lei Zhao was with Stanford University when this work was done. He is now with Jump Operations LLC. Email: zhaolei122@alumni.stanford.edu

Yeow-Khiang Chia was with Stanford University when this work was done. He is now with Institute for Infocomm Research, Singapore. Email: yeowkhiang@gmail.com

Tsachy Weissman is with Stanford University. Email: tsachy@stanford.edu

Material presented in part at Allerton Conference on Communications, Control and Computing, 2011.

tion III, where we characterize the rate-cost tradeoff function for the setting in Fig. 2. We also characterize the rate-cost function when S^n is only causally known at the action encoder. In section IV, we extend the setting to the lossy case where the decoder requires a lossy version of Y^n . We characterize the rate-distortion cost function when S^n is causally known at the action encoder and the side information Z^n is available at both the compressor and the decoder. For other settings, we give achievable schemes for the rate-distortion cost functions. We conclude in Section V, where we mention some possible extensions for future consideration.

II. DEFINITIONS

We give formal definitions for the setups under consideration in this section. We will follow the notation of [11]. Sources (S^n, Z^n) are assumed to be i.i.d.; i.e. $(S^n, Z^n) \sim \prod_{i=1}^n p_{S,Z}(s_i, z_i)$.

A. Lossless case with no side information at the compressor

We now give the definitions for the case when the side information Z^n is not available at the compressor. Referring to Figure 2, a $(n, 2^{nR})$ code for this setup consists of

- an action encoding function $f_a : S^n \rightarrow \mathcal{A}^n$;
- a compression function $f_c : \mathcal{Y}^n \rightarrow M \in [1 : 2^{nR}]$;
- a decoding function $f_d : [1 : 2^{nR}] \times \mathcal{Z}^n \rightarrow \hat{Y}^n$.

The average cost of the system is $E\Lambda(A^n, S^n, Y^n) \triangleq \frac{1}{n} \sum_{i=1}^n E\Lambda(A_i, S_i, Y_i)$. A rate-cost tuple (R, B) is said to be achievable if there exists a sequence of codes such that

$$\limsup_{n \rightarrow \infty} \Pr(Y^n \neq f_d(f_c(Y^n), Z^n)) = 0, \quad (1)$$

$$\limsup_{n \rightarrow \infty} E\Lambda(A^n, S^n, Y^n) \leq B, \quad (2)$$

where $\Lambda(A^n, S^n, Y^n) = \sum_{i=1}^n \Lambda(A_i, S_i, Y_i)/n$. Given cost B , the *rate-cost function*, $R(B)$, is then the infimum of rates R such that (R, B) is achievable.

BC

Remark: Suppose that the channel is given by $P_{Y|A,S} = 1_{Y=A}$ (where $1_{(\cdot)}$ is the indicator function) and that the cost constraint is given by $\Lambda(A^n, S^n)$, then we recover the standard lossy source coding setting with A^n being the reconstruction sequence.

EC

B. Lossless case when side information is available at the compressor

In the case when side information Z^n is available at the compressor, the definitions remain mostly the same, with the exception that the compression function is now given by

$$f_c : \mathcal{Y}^n \times \mathcal{Z}^n \rightarrow M \in [1 : 2^{nR}].$$

C. Lossy case

In the setting where the decoder requires a lossy version of Y^n , the definitions remain largely the same, with the exception that the probability of error constraint, inequality (1), is replaced by the following distortion constraint.

$$\limsup_{n \rightarrow \infty} E d(Y^n, \hat{Y}^n) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_i^n E d(Y_i, \hat{Y}_i) \leq D. \quad (3)$$

A rate R is said to be achievable if there exists a sequence of $(n, 2^{nR})$ codes satisfying both the cost constraint (inequality 2) and the distortion constraint (inequality 3). Given cost B and distortion D , the *rate-cost-distortion function*, $R(B, D)$, is then the infimum of rates R such that the tuple (R, B, D) is achievable.

D. Causal observations of state sequence

In both the lossless and lossy case, we will also consider the setup when the state sequence is only causally known at the action encoder. The definitions remain the same, except for the action encoding function which is now restricted to the following form: For each $i \in [1 : n]$, $f_{a,i} : \mathcal{S}^i \rightarrow \mathcal{A}$.

III. LOSSLESS CASE

In this section, we present our main results for the lossless case. For the lossless case, we will only consider the case when the side information is not available at the compressor, as it will be clear from the results that the presence of side information at the compressor does not change the rate-cost regions for both the case when S^n is causally known, and the case when S^n is non-causally known.

Theorem 1 gives the rate-cost function when the state sequence is non-causally available at the action encoder, while Theorem 2 gives the rate-cost function when the state sequence is causally available.

A. Lossless, non-causal compression with action

Theorem 1 (Rate-cost function for lossless, non-causal case)

The rate-cost function for the compression with action setup when state sequence S^n is non-causally available at the action encoder is given by

$$R(B) = \min_{p(v|s), a=f(s,v): E\Lambda(S,A,Y) \leq B} I(V; S|Z) + H(Y|V, Z), \quad (4)$$

where the joint distribution is of the form $p(z, s, v, a, y) = p(z, s)p(v|s)1_{\{f(s,v)=a\}}p(y|a, s)$. The cardinality of the auxiliary random variable V is upper bounded by $|\mathcal{V}| \leq |\mathcal{S}| + 2$.

Remarks

- Replacing $a = f(s, v)$ by a general distribution $p(a|s, v)$ does not decrease the minimum in (4). For any joint distribution $p(s)p(s|v)p(a|s, v)$, we can always find a random variable W and a function f such that W is independent of S, V and Y , and $A = f(V, W, X)$. Consider $V' = (V, W)$. The Markov condition $V' - (A, S) - (Y, Z)$

still holds. Thus $H(Y|V', Z) + I(V'; S|Z)$ is achievable. Furthermore,

$$\begin{aligned} I(V'; S|Z) + H(Y|V', Z) \\ &= I(V, W; S|Z) + H(Y|V, W, Z) \\ &\leq I(V, W; S|Z) + H(Y|V, Z) \\ &= I(V; S|Z) + H(Y|V, Z). \end{aligned}$$

- $R(B)$ is a convex function in B .
- For each cost function $\Lambda(s, a, y)$, we can replace it with a new cost function involving only s and a by defining $\Lambda'(s, a) = E[\Lambda(S, A, Y)|S = s, A = a]$. Note that Y is distributed as $p(y|s, a)$ given $S = s, A = a$.

BC

- If we set $P_{Y|A, S} = 1_{Y=A}$ and $Z = \emptyset$, then we have

$$\begin{aligned} I(V; S) + H(A|V) &= I(V, A; S) - I(A; S|V) + H(A|V) \\ &= I(V, A; S) \\ &\geq I(A; S). \end{aligned}$$

The rate-cost function then works out to

$$R(B) \geq \min_{\Lambda(A, S) \leq B} I(A; S)$$

for some $p(a|s)$. This recovers the standard lossy source coding result with A being the reconstruction alphabet and B being the desired distortion.

EC

Achievability of Theorem 1 involves an interesting observation in the decoding operation, but before proving the theorem, we first state a corollary of Theorem 1, the case when side information is absent ($Z = \emptyset$). We will also sketch an alternative achievability proof for the corollary, which will serve as a contrast to the achievability scheme for Theorem 1.

Corollary 1 (Side information is absent) *If $Z = \emptyset$, then rate-cost function is given by*

$$R(B) = \min_{p(v|s), a=f(s, v): E\Lambda(S, A, Y) \leq B} I(V; S) + H(Y|V)$$

for some $p(s, v, a, y) = p(s)p(v|s)1_{\{f(s, v)=a\}}p(y|a, s)$.

Achievability sketch for Corollary 1

Code book generation: Fix $p(v|s)$ and $f(s, v)$ and $\epsilon > 0$.

- Generate $2^{n(I(V; S)+\epsilon)}$ $v^n(l)$ sequences independently, $l \in [1 : 2^{n(I(V; S)+\epsilon)}]$, each according to $\prod p_V(v_i)$ to cover S^n .
- For each V^n sequence, the Y^n sequences that are jointly typical with V^n are indexed by $2^{n(H(Y|V)+\epsilon)}$ numbers.

Encoding and Decoding:

- The action encoder looks for a V^n in the code book that is jointly typical with S^n and generates $A_i = f(S_i, V_i)$, $i = 1, \dots, n$.
- The compressor looks for a \hat{V}^n in the codebook that is jointly typical with the channel output Y^n and sends the index of that \hat{V}^n sequence to the decoder. The compressor then sends the index of Y^n as described in the second part of code book generation.

- The decoder simply uses both indices from the compressor to reconstruct Y^n .

Using standard typicality arguments, we can show that the encoding succeeds with high probability and the probability of error can be made arbitrarily small.

Remark: Note that the index of \hat{V}^n is not necessarily equal to V^n . That is, the V^n codeword chosen by the action encoder can be different from the \hat{V}^n codeword chosen by the compressor. But this is not an error event since we still recover the same Y^n even if a different V^n codeword was used.

This scheme, however, does not extend to the case when side information is available at the decoder. The term $H(S|Z, V)$ in Theorem 1 requires us to bin the set of Y^n sequences according to the side information available at the decoder. If we were to extend the above achievability scheme, we would bin the set of Y^n sequences to $2^{n(H(Y|Z, V)+\epsilon)}$ bins. The compressor would find a \hat{V}^n sequence that is jointly typical with Y^n , send the index to the decoder using a rate of $I(V; S|Z) + \epsilon$, and then send the index of the bin which contains Y^n . The decoder would then look for the unique Y^n sequence in the bin that is jointly typical with \hat{V}^n and Z^n . Unfortunately, while the \hat{V}^n codeword is jointly typical with Y^n with high probability, it is not necessarily jointly typical with Z^n , since \hat{V}^n may not be equal to V^n (V^n is jointly typical with Z^n with high probability as V^n is jointly typical with S^n with high probability and $V - S - Z$). One could try to overcome this problem by insisting that the compressor finds the same V^n sequence as the action encoder, but this requirement imposes additional constraints on the achievable rate.

Instead of requiring that the compressor finds a jointly typical V^n sequence, we use an alternative approach to prove Theorem 1. We simply bin the set of all Y^n sequences to $2^{n(I(V; S|Z)+H(Y|Z, V)+\epsilon)}$ bins and send the bin index to the decoder. The decoder looks for the *unique* Y^n sequence in bin M such that $(V^n(l), Y^n, Z^n)$ are jointly typical for *some* $l \in [1 : 2^{n(I(V; S)+\epsilon)}]$. Note that there can more than one $V^n(l)$ sequence which is jointly typical with (Y^n, Z^n) , but this is not an error event as long as the Y^n sequence in bin M is unique. We now give the details of this achievability scheme.

Proof of achievability for Theorem 1

Codebook generation

- Generate $2^{n(I(V; S)+\epsilon)}$ V^n codewords according to $\prod_{i=1}^n p(v_i)$
- For the entire set of possible Y^n sequences, bin them uniformly at random to 2^{nR} bins, where $R > I(V; S) - I(V; Z) + H(Y|Z, V)$, $\mathcal{B}(M)$.

Encoding

- Given s^n , the encoder looks for a v^n sequence in the codebook such that $(v^n, s^n) \in \mathcal{T}_\epsilon^{(n)}$. If there is more than one, it randomly picks one from the set of typical sequences. If there is none, it picks a random index from $[1 : 2^{n(I(V; S)+\epsilon)}]$.

- It then generates a^n according to $a_i = f(v_i, s_i)$ for $i \in [1 : n]$.
- At the second encoder, it takes the output y^n sequences and sends out the bin index M such that $y^n \in \mathcal{B}(M)$.

Decoding

- The decoder looks for the *unique* \hat{y}^n sequence such that $(v^n(l), \hat{y}^n, z^n) \in \mathcal{T}_\epsilon^{(n)}$ for *some* $l \in [1 : 2^{n(I(V;S))}]$ and $\hat{y}^n \in \mathcal{B}(M)$. If there is none or more than one, it declares an error.

Analysis of probability of error

Define the following error events

$$\begin{aligned}\mathcal{E}_0 &:= \{(V^n(L), Z^n, Y^n) \notin \mathcal{T}_\epsilon^{(n)}\}, \\ \mathcal{E}_l &:= \{(V^n(l), Z^n, \hat{Y}^n) \in \mathcal{T}_\epsilon^{(n)} \\ &\quad \text{for some } \hat{Y}^n \neq Y^n, \hat{Y}^n \in \mathcal{B}(M)\}.\end{aligned}$$

By symmetry of the codebook generation, it suffices to consider $M = 1$. The probability of error is upper bounded by

$$P(\mathcal{E}) \leq P(\mathcal{E}_0) + \sum_{l=1}^{2^{n(I(V;S)) + (\epsilon)}} P(\mathcal{E}_l).$$

$P(\mathcal{E}_0) \rightarrow 0$ as $n \rightarrow \infty$ following standard analysis of probability of error. It remains to analyze the second error term. Consider $P(\mathcal{E}_l)$ and define $\mathcal{E}_l(V^n, Z^n) := \{(V^n(l), Z^n, \hat{Y}^n) \in \mathcal{T}_\epsilon^{(n)} \text{ for some } \hat{Y}^n \neq Y^n, \hat{Y}^n \in \mathcal{B}(1)\}$. We have

$$\begin{aligned}P(\mathcal{E}_l) &= P(\mathcal{E}_l(V^n, Z^n)) \\ &= \sum_{(v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}} P(V^n(l) = v^n, Z^n = z^n) P(\mathcal{E}_l(v^n, z^n) | v^n, z^n) \stackrel{(b)}{=} \sum_{(v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}} P(V^n(l) = v^n, Z^n = z^n) \\ &\quad \left(\sum_{y^n} P(Y^n = y^n | v^n, z^n) P(\mathcal{E}_l(v^n, z^n) | v^n, z^n, y^n) \right) \\ &\stackrel{(a)}{\leq} \sum_{(v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}} (P(V^n(l) = v^n, Z^n = z^n) \\ &\quad \sum_{y^n} P(Y^n = y^n | v^n, z^n) 2^{n(H(Y|Z,V) + (\epsilon) - R)}) \\ &\stackrel{(b)}{=} \sum_{(v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}} (P(V^n(l) = v^n) P(Z^n = z^n) \\ &\quad 2^{n(H(Y|Z,V) + (\epsilon) - R)}) \\ &\leq \left(2^{n(H(V,Z) + (\epsilon))} 2^{-n(H(V) - (\epsilon))} 2^{-n(H(Z) - (\epsilon))} \right. \\ &\quad \left. 2^{n(H(Y|Z,V) + (\epsilon) - R)} \right) \\ &= 2^{n(H(Y|V,Z) - I(V;Z) - R - 4(\epsilon))}.\end{aligned}$$

(a) follows since the set of Y^n sequences are binned uniformly at random independent of other Y^n sequences, and

the fact that there are at most $2^{n(H(Y|Z,V) + (\epsilon))}$ Y^n sequences which are jointly typical with a given typical (v^n, z^n) . (b) follows from the fact that the codebook generation is independent of (S^n, Z^n) . Therefore, for any fixed l , $V^n(l)$ is independent of Z^n . Hence, if $R \geq I(V;S) - I(V;Z) + H(Y|Z,V) + 6(\epsilon)$,

$$\sum_{l=1}^{2^{n(I(V;S) + (\epsilon))}} P(\mathcal{E}_l) \leq 2^{-n(\epsilon)} \rightarrow 0,$$

as $n \rightarrow \infty$.

We now turn to the proof of converse for Theorem 1

Proof of converse for Theorem 1

Given a $(n, 2^{nR})$ code for which the probability of error goes to zero with n and satisfies the cost constraint, define $V_i = (Z^{n \setminus i}, S_{i+1}^n, Y^{i-1})$, we have

$$\begin{aligned}nR &\geq H(M|Z^n) \\ &= H(M, Y^n|Z^n) - H(Y^n|M, Z^n) \\ &\stackrel{(a)}{=} H(M, Y^n|Z^n) - n\epsilon_n \\ &= H(Y^n|Z^n) - n\epsilon_n \\ &= \sum_{i=1}^n H(Y_i|Y^{i-1}, Z^n) - n\epsilon_n \\ &= \sum_{i=1}^n H(Y_i|Y^{i-1}, S_{i+1}^n, Z^n) \\ &\quad + \sum_{i=1}^n I(Y_i; S_{i+1}^n | Y^{i-1}, Z^n) - n\epsilon_n \\ &\stackrel{(b)}{=} \sum_{i=1}^n H(Y_i|Y^{i-1}, S_{i+1}^n, Z^n) \\ &\quad + \sum_{i=1}^n I(Y^{i-1}; S_i | S_{i+1}^n, Z^n) - n\epsilon_n \\ &\stackrel{(c)}{=} \sum_{i=1}^n H(Y_i|Y^{i-1}, S_{i+1}^n, Z^n) \\ &\quad + \sum_{i=1}^n I(Y^{i-1}, S_{i+1}^n, Z^{n \setminus i}; S_i | Z_i) - n\epsilon_n \\ &\stackrel{(d)}{=} \sum_{i=1}^n H(Y_i|V_i, Z^i) + \sum_{i=1}^n I(V_i; S_i | Z^i) - n\epsilon_n \\ &= nH(Y_Q, |V_Q, Q, Z_Q) + nI(V_Q; S_Q | Q, Z_Q) - n\epsilon_n\end{aligned}$$

where (a) is due to Fano's inequality. (b) follows from Csiszár Sum. (c) holds because (S^n, Z^n) is an i.i.d source. Note that the Markov conditions, $V_i - (S_i, A_i) - Y_i$ and $V_i - S_i - Z_i$ hold. Finally, we introduce Q as the time sharing random variable, i.e., $Q \sim \text{Unif}[1, \dots, n]$, and set $V = (V_Q, Q)$, $Y = Y_Q$ and $S = S_Q$, which completes the proof.

Remark: Note that the proof of converse continues to hold even if side information Z^n is available at the compressor. This observation shows that side information Z^n at the compressor does not change the rate-cost tradeoff region in Theorem 1.

B. Lossless, causal compression with action

Our next result gives the rate-cost function for the case of lossless, causal compression with action.

Theorem 2 (Rate-cost function for lossless, causal case)

The rate for the compression with action when the state information is causally available at the action encoder is given by

$$R(B) = \min_{p(v), a=f(s,v): EA(S,A,Y) \leq B} H(Y|V, Z) \quad (5)$$

where the joint distribution is of the form $p(z, s, v, a, y) = p(z, s)p(v)1_{\{f(s,v)=a\}}p(y|a, s)$. The cardinality of V is upper bounded by $|S| + 2$.

Achievability sketch: Here V simply serves as a time-sharing random variable. Fix a $p(v)$ and $f(s, v)$. We first generate a V^n sequence and reveal it to the action encoder, the compressor and the decoder. The encoder generates $A_i = f(S_i, V_i)$. The compressor simply bins the set of Y^n sequences to $2^{n(H(Y|V, Z) + \epsilon)}$ bins and sends the index of the bin which contains Y^n . The decoder recovers Y^n by finding the unique Y^n sequence in bin M such that (V^n, Z^n, Y^n) are jointly typical.

Remark: Just as in the non-causal case, the achievability is closely related to the channel coding strategy in [2], our achievability in this section uses the ‘‘Shannon Strategy’’ in [4]. In both cases, the optimal channel coding strategy yield the most compressible output when the message rate goes to zero.

Proof of Converse: Given a $(n, 2^{nR})$ code that satisfies the constraints, define $V_i = (S^{i-1}, Z^{n \setminus i})$. We have

$$\begin{aligned} nR &\geq H(M|Z^n) \\ &= H(M, Y^n|Z^n) - H(Y^n|M, Z^n) \\ &\stackrel{(a)}{=} H(M, Y^n|Z^n) - n\epsilon_n \\ &= H(Y^n|Z^n) - n\epsilon_n \\ &= \sum_{i=1}^n H(Y_i|Y^{i-1}, Z_i, Z^{n \setminus i}) - n\epsilon_n \\ &\geq \sum_{i=1}^n H(Y_i|Y^{i-1}, A^{i-1}, S^{i-1}, Z_i, Z^{n \setminus i}) - n\epsilon_n \\ &\stackrel{(b)}{=} \sum_{i=1}^n H(Y_i|A^{i-1}, S^{i-1}, Z_i, Z^{n \setminus i}) - n\epsilon_n \\ &\stackrel{(c)}{=} \sum_{i=1}^n H(Y_i|V_i, Z_i) - n\epsilon_n \\ &\stackrel{(d)}{=} nH(Y_Q|V_Q, Q, Z_Q) - n\epsilon_n \end{aligned}$$

where (a) is due to Fano’s inequality; (b) follows from the Markov chain $Y_i - (S^{i-1}, A^{i-1}, Z^n) - Y^{i-1}$; (c) follows since A^{i-1} is a function of S^{i-1} . Note that A_i is now a function of S_i and V_i . Finally, we introduce Q as the time sharing random variable, i.e., $Q \sim \text{Unif}[1, \dots, n]$. Thus, by setting $V = (V_Q, Q)$ and $Y = Y_Q$, we have completed the proof.

Remark: Note that the proof of converse continues to hold even if side information Z^n is available at the compressor. This observation shows that side information Z^n at the

compressor does not change the rate-cost tradeoff region in Theorem 2.

C. Examples

1) *No side information:* In this subsection, we first consider an example with state sequence $S^n \sim \text{i.i.d. Bern}(1/2)$ and $Z = \emptyset$. We have two actions available, $A = 0$ and $A = 1$. The cost constraint is on the frequency of action $A = 1$, $EA \leq B$. The channel output $Y_i = S_i \oplus A_i \oplus S_{Ni}$ where \oplus is the modulo 2 sum and $\{S_{Ni}\}$ are i.i.d. $\text{Bern}(p)$ noise, $p < 1/2$. The example is illustrated in Fig. 3.




Fig. 3. Binary example with side information $Z = \emptyset$.

We use the following lemma to simplify the optimization problem in Eq. (4) applied to the binary example.

Lemma 1 For the binary example, it is without loss of optimality to have the following constraints when solving the optimization problem of Eq. (4):

- $\mathcal{V} = \{0, 1, 2\}$, $\Pr(V = 0) = \Pr(V = 1) = \theta/2$, for some $\theta \in [0, 1]$.
- The function $a = f(s, v)$ is of the form: $f(s, 0) = s$, $f(s, 1) = 1 - s$ and $f(s, 2) = 0$.
- $\Pr(S = 0|V = 1) = \Pr(S = 1|V = 0) = \Delta$ and $\Pr(S = 0|V = 2) = 1/2$.
- $\Delta\theta \leq B$.

Note that the constraints guarantee that $\Pr(S = 0) = \Pr(S = 1) = 1/2$.

Proof: See Appendix. ■

Using Lemma 1, we can simplify the objective function in Eq. (4) in the following way:

$$\begin{aligned} &H(Y|V) + I(V; S) \\ &= H(Y|V) - H(S|V) + H(S) \\ &= H(S \oplus A \oplus S_N|V) - H(S|V) + 1 \\ &= \frac{\theta}{2} (H(0 \oplus S_N|V = 0) - H(\Delta)) \\ &\quad + \frac{\theta}{2} \{H(1 \oplus S_N|V = 1) - H(\Delta)\} \\ &\quad + (1 - \theta) \{H(S \oplus S_N|V = 2) - 1\} + 1 \\ &= \theta(H_2(p) - H(\Delta)) + 1 \end{aligned}$$

where $H_2(\cdot)$ is the binary entropy function, i.e., $H_2(\delta) = -\delta \log \delta - (1 - \delta) \log(1 - \delta)$.

$$\begin{aligned}
R(B) &= \min_{\theta \in [2B, 1], \Delta \leq B} \theta (H_2(p) - H(\Delta)) + 1 \\
&= 1 + \min_{\Delta \in [B, 1/2]} \frac{B}{\Delta} (H_2(p) - H_2(\Delta)) \\
&= 1 - B \max_{\Delta \in [B, 1/2]} \frac{H_2(\Delta) - H_2(p)}{\Delta} \\
&= \begin{cases} 1 - B \frac{H_2(b^*) - H_2(p)}{b^*}, & \text{if } 0 \leq B < b^* \\ 1 - H_2(B) + H_2(p), & \text{if } b^* \leq B \leq 1/2 \end{cases} \quad (6)
\end{aligned}$$

where b^* is the solution of the following function:

$$\frac{H_2(b) - H_2(p)}{b} = \frac{dH_2}{db}, \quad b \in [0, 1/2] \quad (7)$$

which is illustrated in Fig. 4.




Fig. 4. The threshold b^* solves $\frac{H_2(b) - H_2(p)}{b} = \frac{dH_2}{db}$, $b \in [0, 1/2]$

Now let us shift our attention to the causal case of the binary example, i.e., S_i is only causally available at the action encoder.

Lemma 2 *For the causal case of the binary example, it is without loss of optimality to have the following constraints when solving the optimization problem in Eq. (5):*

- $\mathcal{V} = \{0, 1\}$, $\Pr(V = 0) = \theta$, for some $\theta \in [0, 1]$.
- The function $a = f(s, v)$ is of the form: $f(s, 0) = s$, $f(s, 1) = 0$.
- $\frac{\theta}{2} \leq B$.

Proof: See Appendix. ■

$$\begin{aligned}
R(B) &= \min H(Y|V) \\
&= \min_{\theta \in [0, 1], \frac{\theta}{2} \leq B} \theta H(Y|V=0) + (1 - \theta) H(Y|V=1) \\
&= \min_{\theta \in [0, 1], \frac{\theta}{2} \leq B} \theta H(Z|V=0) + (1 - \theta) H(S \oplus Z|V=1) \\
&= \min_{\theta \in [0, 1], \frac{\theta}{2} \leq B} \theta H_2(p) + (1 - \theta) \\
&= \begin{cases} 2B H_2(p) + (1 - 2B), & 0 \leq B \leq 1/2; \\ H_2(p), & 1/2 \leq B. \end{cases}
\end{aligned}$$

For the binary example with $p = 0.1$, we plot the rate-cost function $R(B)$ for both cases in Figure 5. Note that when S is only causally known at the action encoder, the optimum lossless compression scheme amounts to time sharing between compressing the noise S_N losslessly and compressing S losslessly. The optimum time sharing factor is determined by the cost B .




Fig. 5. Comparison between the non-causal and causal rate-cost functions. The parameter of the Bernoulli noise is set at 0.1.

2) *Erased side information:* We now turn to the case when side information is available at the decoder only. We extend our setting in the previous example by letting Z be an erased version of S . That is,

$$Z = \begin{cases} S & \text{w.p. } 1 - p_e \\ e & \text{w.p. } p_e \end{cases}$$

In this case, the rate cost function is related to the case when no side information is available at the decoder in a simple manner. We first note the following

$$\begin{aligned}
P(V|Z=e) &= \sum_S P(S, V|Z=e) \\
&= \sum_S P(S|Z=e) P(V|S, Z=e) \\
&= \sum_S P(S) P(V|S) \\
&= P(V). \quad (8)
\end{aligned}$$

The third equality follows from the Markov Chain $V - S - Z$. Furthermore,

$$\begin{aligned}
P(Y|V, Z=e) &= \sum_S p(S|Z=e, V) p(Y|V, S, Z=e) \\
&= \sum_S P(S|V) P(Y|V, S) \\
&= P(Y|V). \quad (9)
\end{aligned}$$

The second equality follows from the Markov chain $Z - (S, V) - Y$ and $P(S|Z=e, V) = P(S, V, Z=e) / P(V, Z=e) = P(S, V) P(Z=e) / (P(Z=e) P(v)) = P(S|V)$. We

now consider the rate-cost expression when S is non-causally known at the action encoder.

$$\begin{aligned} R(B)_{\text{NC}} &= \min I(V; S|Z) + H(Y|V, Z) \\ &\stackrel{(a)}{=} \min p_e I(V; S) + p_e H(Y|V) + (1 - p_e) H(Y|V, S) \\ &= p_e \min(I(V; S) + H(Y|V)) + (1 - p_e) H_2(p). \end{aligned}$$

(a) follows the following observations: (i) when $Z = \epsilon$, $P(V|Z = \epsilon) = P(V)$ by (8), so $H(V|Z = \epsilon) = H(V)$ and from the Markov Chain $V - S - Z$ and $P(S|Z = \epsilon) = P(S)$, $H(V|S, Z = \epsilon) = H(V|S)$. Hence, $I(V; S|Z = \epsilon) = I(V; S)$; and (ii) from (9), $H(Y|V, Z = \epsilon) = H(Y|V)$. The last equality follows from when $Z = S$, $H(Y|Z = S, V) = H(Y|S, V)$. Since $A = f(S, V)$, $S_N \sim \text{Bern}(p)$ independent of (S, V, Z) and $Y = S \oplus A \oplus S_N$, $H(Y|S, V) = H_2(p)$.

As checks, note that when $p_e = 1$, which corresponds to the no side information case, the rate-cost function reduces to that in Corollary 1, and when $p_e = 0$, the rate-cost function reduces to $H_2(p)$, which corresponds to the minimum rate required when S is also available at the decoder. We now turn to the case when S is only causally known at the action encoder. Here, we have

$$\begin{aligned} R(B)_C &= \min H(Y|V, Z) \\ &= p_e \min H(Y|V) + (1 - p_e) H_2(p). \end{aligned}$$

The rate-cost tradeoff is shown in figure 6.




Fig. 6. Comparison between the non-causal and causal rate-cost functions with erased side information at the decoder. The parameter of the Bernoulli noise is set at 0.1 and the parameter of the erased side information is set at 0.5.

IV. LOSSY COMPRESSION WITH ACTIONS

In this section, we extend our setup to the lossy case. We first consider the case when side information is available at both the compressor and the decoder. We characterize the rate-distortion-cost tradeoff region for the case when S^n is causally known to the action encoder. The case when S^n is non-causally known at the action encoder is more involved. We give an achievable rate-distortion-cost region for that setting.

We then move on to the case when side information Z is available at the decoder only, and S is non-causally known at the action encoder. We discuss two achievability schemes for this setting.

A. Side information known at compressor and decoder

Theorem 3 The rate-cost-distortion function for the case with causal state information and side information available to both the compressor and the decoder is given by

$$R(B, D) = \min_{a=f(s,v): E\Lambda(S,A,Y) \leq B, E d(Y, \hat{Y}) \leq D} I(Y; \hat{Y}|V, Z) \quad (10)$$

where the joint distribution is of the form $p(z, s, v, a, y, \hat{y}) = p(s, z)p(v)1_{\{a=f(s,v)\}}p(y|a, s)p(\hat{y}|y, v, z)$. The cardinality of V is upper bounded by $|S| + 2$.

Achievability sketch: The achievability is straightforward, with V^n acting as the time sharing random variable known to all parties. We first generate $V^n \sim \prod_{i=1}^n p(v_i)$. For each z^n sequence, we generate $2^{n(I(\hat{Y}; Y|V, Z) + \epsilon)}$ \hat{Y}^n sequences according to $\prod_{i=1}^n p(\hat{y}_i|v_i, z_i)$. The action encoder simply generates a^n according to $a_i = f(v_i, s_i)$ for $i \in [1 : n]$. The compressor looks for a \hat{y}^n sequence such that $(\hat{y}^n, y^n, v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}$. It then sends out this description to the decoder which reconstructs Y^n as \hat{y}^n . Since we have $2^{n(I(\hat{Y}; Y|V, Z) + \epsilon)}$ \hat{Y}^n sequences, the probability of not finding a jointly typical \hat{Y}^n sequence goes to zero as $n \rightarrow \infty$.

Converse: Given a $(n, 2^{nR})$ code satisfying the cost and distortion conditions, we have

$$\begin{aligned} nR &\geq H(M|Z^n) \\ &\geq I(M; Y^n|Z^n) \\ &= \sum_{i=1}^n I(M; Y_i|Y^{i-1}, Z^{n \setminus i}, Z^i) \\ &\stackrel{(a)}{=} \sum_{i=1}^n I(M; Y_i|V_i, Z_i) \\ &\stackrel{(b)}{\geq} \sum_{i=1}^n I(\hat{Y}_i; Y_i|V_i, Z_i) \\ &\stackrel{(c)}{=} nI(\hat{Y}_Q; Y_Q|V_Q, Q, Z_Q) \end{aligned}$$

where in (a) we set $V_i = (Y^{i-1}, Z^{n \setminus i})$. (b) holds from the fact that \hat{Y}_i is a function of M and Z^n . In (c) we introduce Q as the time sharing random variable, i.e., $Q \sim \text{Unif}[1, \dots, n]$. Thus, by setting $V = (V_Q, Q)$ and $Y = Y_Q$ and noting that V is independent of (S, Z) , we have shown that $R(B, D) \geq I(\hat{Y}, Y|V, Z)$ for some $p(v)p(s|v)p(z|s)p(a|s, v)p(y|a, s)p(\hat{y}|y, a, s, z, v)$. It suffices to restrict attention to the joint distribution stated in Theorem 3 because of the following observations.

- $p(\hat{y}|a, y, v, s, z)$ can be restricted to $p(\hat{y}|y, v, z)$ since the mutual information term $I(Y; \hat{Y}|V, Z)$ and the distortion constraint only depend on the marginal distribution $p(\hat{y}, y, v, z)$.
- $p(a|s, v)$ can be restricted to $a = f(s, v)$ since we can always find an independent random variable U

such that $p(a|s, v) = p(s, v)p(u)1_{a=f(s, v, u)}$. Now define $\tilde{V} = (V, U)$ and $p(\hat{y}|y, \tilde{v}, z) = p(\hat{y}|y, \tilde{v}, z)$. Note that $p(s, z)p(\tilde{v})1_{\{a=f(s, \tilde{v})\}}p(y|a, s)p(\hat{y}|y, \tilde{v}, z) = p(s, z)p(v)p(a|s, v)p(y|a, s)p(\hat{y}|y, v, z)$. Since the joint distribution remains unchanged, the distortion and cost are preserved. As for the rate, we note that

$$\begin{aligned} I(\hat{Y}; Y|\tilde{V}, Z) &= H(\hat{Y}|\tilde{V}, Z) - H(\hat{Y}|\tilde{V}, Z, Y) \\ &= H(\hat{Y}|V, U, Z) - H(\hat{Y}|V, U, Z, Y) \\ &\leq H(\hat{Y}|V, Z) - H(\hat{Y}|V, Z, Y) \\ &= I(\hat{Y}; Y|V, Z). \end{aligned}$$

Our next Theorem gives an upper bound on the rate-distortion-cost tradeoff for the case when the state information is known non-causally at the action encoder and Z^n is present at both the compressor and the decoder.

Theorem 4 *An upper bound on the rate-distortion-cost function for the case with non-causal state information and side information at both the compressor and decoder is given by*

$$R(B, D) \leq \min_{E\Lambda(S, A, Y) \leq B, Ed(Y, \hat{Y}) \leq D} I(V; S|Z) + I(\hat{Y}; Y|V, Z) \quad (11)$$

where the joint distribution is of the form $p(s, v, a, y, \hat{y}, z) = p(s, z)p(v|s)1_{\{f(s, v)=a\}}p(y|a, s)p(\hat{y}|y, v, z)$.

Sketch of achievability:

We generate $2^{n(I(V; S)+\epsilon)}$ $V^n(l_0)$, $l_0 \in [1 : 2^{n(I(V; S)+\epsilon)}]$, sequences according to $\prod_{i=1}^n p(v_i)$, and for each $v^n(l_0)$ and z^n , generate $2^{n(I(\hat{Y}; Y|V, Z)+\epsilon)}$ $\hat{Y}^n(l_0, l_1)$, $l_1 \in [1 : 2^{n(I(\hat{Y}; Y|V, Z)+\epsilon)}]$, sequences according to $\prod_{i=1}^n p(\hat{y}_i|v_i, z_i)$. The set of v^n sequences are then randomly binned to $2^{n(I(V; S|Z)+2\epsilon)}$ bins, $\mathcal{B}(m)$, $m \in [1 : 2^{n(I(V; S|Z)+2\epsilon)}]$. Given a sequence s^n , the action encoder finds the v^n sequence which is jointly typical with s^n and takes actions according to $a_i = f(s_i, v_i)$ for $i \in [1 : n]$. At the compressor, we first find a $\tilde{v}^n(l_0)$ that is jointly typical with (y^n, z^n) and then, a $\hat{y}^n(l_0, l_1)$ such that $(\tilde{v}^n(l_0), \hat{y}^n(l_0, l_1), y^n, z^n) \in \mathcal{T}_\epsilon^{(n)}$. Note that there exists at least one $\tilde{v}^n(l_0)$ that is jointly typical with (y^n, z^n) with high probability since the true v^n sequence is jointly typical with (y^n, z^n) with high probability. If there is more than one such sequence, the compressor chooses one uniformly at random from the set of \tilde{v}^n sequences jointly typical with (y^n, z^n) . The compressor then sends the indices m and l_1 such that the selected $\tilde{v}^n(l_0) \in \mathcal{B}(m)$. The decoder recovers $\tilde{v}^n(l_0)$ by looking for the unique $\hat{l}_0 \in \mathcal{B}(m)$ such that $(\hat{v}^n(\hat{l}_0), z^n) \in \mathcal{T}_\epsilon^{(n)}$. It reconstructs Y^n as $\hat{y}(\hat{l}_0, l_1)$. From the rates given, it is easy to see that all encoding and decoding steps succeed with high probability as $n \rightarrow \infty$.

B. Side information available at the decoder only

When the side information is available at the decoder only and S is known non-causally at the action encoder, we discuss two possible achievability schemes. The first scheme is a generalization of the achievability scheme of Theorem 1 to the lossy case.

Theorem 5 *An upper bound on the rate-distortion-cost function for the case with non-causal state information and side information at the decoder is given by*

$$R(B, D) = I(V; S|Z) + I(U; Y|V, Z)$$

for some $p(z, s)p(v|s)1_{a=f(s, v)}p(y|a, s)p(u|y)$ satisfying

$$\begin{aligned} E\Lambda(A, S, Y) &\leq B, \\ E d(Y, \hat{Y}(Z, U)) &\leq D. \end{aligned}$$

Theorem 5 is generalization of Theorem 1 since if we let $U = Y$, we recover Theorem 1.

Proof: As the achievability scheme is an extension of the achievability scheme for Theorem 1, we will only mention the additional steps in the proof of achievability.

Codebook Generation

The additional step in the codebook generation procedure is in the generation of a codebook of U^n covering sequences to cover Y^n and a binning or compression codebook for the U^n sequences. We first generate $2^{n(I(U; Y)+\epsilon)}$ $U^n(l_1)$ sequences according to $\prod_{i=1}^n p(u_i)$. We then bin the set of U^n sequences into $2^{n(I(V; S|Z)+I(U; Y|V, Z)+5\epsilon)}$ bins, $\mathcal{B}(M)$, $M \in [1 : 2^{n(I(V; S|Z)+I(U; Y|V, Z)+5\epsilon)}]$.

Encoding

The encoding procedure for the action encoder remains the same as that in Theorem 1. For the compressor, it first looks for a $u^n(l_1)$ such that $(u^n(l_1), y^n) \in \mathcal{T}_\epsilon^{(n)}$. It then sends out the index m , such that $u^n(l_1) \in \mathcal{B}(m)$.

Decoding and analysis of probability of error

The decoder looks for the unique $u^n(\hat{l}_1) \in \mathcal{B}(m)$ such that $(v^n(l_0), y^n, z^n, u^n(l_1)) \in \mathcal{T}_\epsilon^{(n)}$ for some $l_0 \in [1 : 2^{n(I(V; S)+\epsilon)}]$. For the analysis of probability of error, let L_0 and L_1 be the indices picked by the action encoder and the compressor respectively. Following the rates given in the codebook generation and encoding procedure, the covering lemma [11, Chapter 3] and the strong Markov lemma [11, Chapter 12], it is easy to see that $P(V^n(L_0), Y^n, Z^n, U^n(L_1) \in \mathcal{T}_\epsilon^{(n)}) \rightarrow 1$ as $n \rightarrow \infty$. The other “error” event of interest is now the following.

$$\begin{aligned} \mathcal{E}_U &:= \{(V^n(l_0), Z^n, U^n(\hat{l}_1)) \in \mathcal{T}_\epsilon^{(n)} \\ &\quad \text{for some } \hat{l}_1 \neq L_1, U^n(\hat{l}_1) \in \mathcal{B}(M), \\ &\quad l_0 \in [1 : 2^{n(I(V; S)+\epsilon)}]\}. \end{aligned}$$

Due to the symmetry of the binning process, we can assume without loss of generality, $M = 1$. Define $\mathcal{E}_{l_0}(V^n, Z^n)$ to be the event

$$\begin{aligned} \mathcal{E}_{l_0}(V^n, Z^n) &:= \{(V^n(l_0), Z^n, U^n(\hat{l}_1)) \in \mathcal{T}_\epsilon^{(n)} \\ &\quad \text{for some } \hat{l}_1 \neq L_1, U^n(\hat{l}_1) \in \mathcal{B}(1)\}. \end{aligned}$$

Then, $P(\mathcal{E}_U)$ is upper bounded by

$$P(\mathcal{E}_U) \leq \sum_{l_0} P(\mathcal{E}_{l_0}(V^n, Z^n)). \quad (12)$$

We now give a bound for $P(\mathcal{E}_{l_0}(V^n, Z^n))$. We first have

$$\begin{aligned} & P(\mathcal{E}_{l_0}(V^n, Z^n)) \\ &= \sum_{(v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}} \left(\frac{p_{V,Z}(v^n, z^n)}{P(\mathcal{E}_{l_0}(V^n, Z^n) | V^n = v^n, Z^n = z^n)} \right) \\ &= \sum_{(v^n, z^n) \in \mathcal{T}_\epsilon^{(n)}} \left(\frac{p_{V,Z}(v^n, z^n)}{P(\mathcal{E}_{l_0}(v^n, z^n) | v^n, z^n)} \right). \end{aligned}$$

Next, let $\mathcal{E}_{l_0}(j, v^n, z^n)$ be the error event $\{(v^n(l_0), z^n, U^n(j)) \in \mathcal{T}_\epsilon^{(n)}, L_1 \neq j, U^n(j) \in \mathcal{B}(1)\}$. Then, $\mathcal{E}_{l_0}(v^n, z^n) \subseteq \bigcup_{j=1}^{2^{n(I(U;Y)+\epsilon)}} \mathcal{E}_{l_0}(j, v^n, z^n)$. We therefore have

$$\begin{aligned} & P(\mathcal{E}_{l_0}(v^n, z^n | v^n, z^n)) \\ & \leq \sum_{j=1}^{2^{n(I(U;Y)+\epsilon)}} P(\mathcal{E}_{l_0}(j, v^n, z^n) | v^n, z^n) \\ & \leq \sum_{j=1}^{2^{n(I(U;Y)+\epsilon)}} P((v^n, z^n, U^n(j)) \in \mathcal{T}_\epsilon^{(n)}, U^n(j) \in \mathcal{B}(1) | v^n, z^n) \\ &= \sum_{j=1}^{2^{n(I(U;Y)+\epsilon)}} \left(\frac{P((v^n, z^n, U^n(j)) \in \mathcal{T}_\epsilon^{(n)} | v^n, z^n)}{P(U^n(j) \in \mathcal{B}(1) | v^n, z^n)} \right) \\ & \leq 2^{n(I(U;Y)+\epsilon)} \cdot 2^{-n(I(U;Z,V)-\epsilon)} \cdot 2^{-n(I(V;S|Z)+I(U;Y|V,Z)+5\epsilon)} \\ &= 2^{-n(I(V;S|Z)+3\epsilon)}. \end{aligned}$$

We therefore have

$$\begin{aligned} & P(\mathcal{E}_{l_0}(V^n, Z^n)) \\ & \leq \sum_{v^n, z^n \in \mathcal{T}_\epsilon^{(n)}} P(V^n(l_0) = v^n, Z^n = z^n) \cdot 2^{-n(I(V;S|Z)+3\epsilon)} \\ & \leq 2^{-n(I(V;Z)-\epsilon)} \cdot 2^{-n(I(V;S|Z)+3\epsilon)}. \end{aligned}$$

Hence, from (12),

$$\begin{aligned} P(E_U) & \leq 2^{n(I(V;S)+\epsilon)} 2^{-n(I(V;Z)-\epsilon)} \cdot 2^{-n(I(V;S|Z)+3\epsilon)} \\ &= 2^{-n\epsilon}. \end{aligned}$$

Therefore, $P(E_U) \rightarrow 0$ as $n \rightarrow \infty$.

Since the probability of “error” goes to zero as $n \rightarrow \infty$, the expected distortion of the reconstruction $\hat{Y}_i = \hat{y}(z_i, u_i)$, $i \in [1 : n]$, is less than or equal to D as $n \rightarrow \infty$. ■

The achievability scheme in Theorem 5 restricts the description of Y^n that is sent, U^n , to be independent of V^n given Y^n . This is a result of not requiring the compressor to decode the true V^n codeword that was selected by the action encoder. In our next scheme, we remove the Markov condition, $U - Y - V$, by making the compressor decode V^n . This operation results in a different restriction on the allowable joint probability distribution, $I(V; Y) \geq I(V; S)$.

Theorem 6 *An upper bound on the rate-distortion-cost function for the case with non-causal state information and side information at the decoder is given by*

$$R(B, D) = I(V; S|Z) + I(U; Y|V, Z)$$

for some $p(z, s)p(v|s)1_{a=f(s,v)}p(y|a, s)p(u|y, v)$ satisfying

$$\begin{aligned} E \Lambda(A, S, Y) & \leq B, \\ E d(Y, \hat{Y}(Z, U)) & \leq D, \\ I(V; S) & \leq I(V; Y). \end{aligned}$$

Sketch of achievability

We generate $2^{n(I(V;S)+\epsilon)}$ $V^n(l_0)$, $l_0 \in [1 : 2^{n(I(V;S)+\epsilon)}]$, sequences according to $\prod_{i=1}^n p(v_i)$, and for each $v^n(l_0)$, generate $2^{n(I(U;Y|V)+\epsilon)}$ $U^n(l_0, l_1)$, $l_1 \in [1 : 2^{n(I(U;Y|V)+\epsilon)}]$, sequences according to $\prod_{i=1}^n p(u_i|v_i)$. The set of V^n sequences are partitioned to $2^{n(I(V;S|Z)+2\epsilon)}$ bins, $\mathcal{B}(m_0)$, $m_0 \in [1 : 2^{n(I(V;S|Z)+2\epsilon)}]$, while the set of U^n sequences are partitioned to $2^{n(I(U;Y|V,Z)+2\epsilon)}$ bins, $\mathcal{B}(m_1)$, $m_1 \in [1 : 2^{n(I(U;Y|V,Z)+2\epsilon)}]$. Given a sequence s^n , the action encoder finds the v^n sequence which is jointly typical with s^n and takes actions according to $a_i = f(s_i, v_i)$ for $i \in [1 : n]$. At the compressor, we first find $v^n(l_0)$ by joint typicality decoding. It can be shown that this decoding procedure succeeds with high probability provided $I(V; Y) \geq I(V; S)$ [12]. Next, the compressor looks for a $u^n(l_0, l_1)$ such that $(v^n(l_0), u^n(l_0, l_1), y^n, z^n) \in \mathcal{T}_\epsilon^{(n)}$. The compressor then sends the indices m_0 and m_1 such that $v^n(l_0) \in \mathcal{B}(m_0)$ and $u^n(l_0, l_1) \in \mathcal{B}(m_1)$. The decoding operation now follows standard Wyner-Ziv decoding. The decoder first recovers $v^n(l_0)$ by looking for the unique $\hat{l}_0 \in \mathcal{B}(m_0)$ such that $(\hat{v}^n(\hat{l}_0), z^n) \in \mathcal{T}_\epsilon^{(n)}$. Next, it recovers $u^n(l_0, l_1)$ by looking for the unique $\hat{u}^n(\hat{l}_0, \hat{l}_1)$ such that $(\hat{v}^n(\hat{l}_0), \hat{u}^n(\hat{l}_0, \hat{l}_1), z^n) \in \mathcal{T}_\epsilon^{(n)}$. It reconstructs Y^n as $\hat{y}_i(\hat{v}_i(\hat{l}_0), \hat{u}_i(\hat{l}_0, \hat{l}_1), z_i)$ for $i \in [1 : n]$. From the rates given, the encoding and decoding steps succeed with high probability as $n \rightarrow \infty$.

V. CONCLUSION AND FUTURE DIRECTIONS

In this paper, we consider a variation of lossy and lossless compression where, instead of compressing the original source, we take actions to modify the source before compression, subject to a cost constraint. In the lossless case, we characterize the rate-cost tradeoff for several different cases, including the cases where side information is available at the decoder only, and where the original source S^n is either known causally or non-causally at the action encoder. We then extended the analysis to the lossy case, where we characterize the rate-distortion-cost tradeoff for the case where S^n is known only causally and side information is available at both the compressor and the decoder.

Our setting can be extended in several different directions. One possible extension is to consider the case of message embedding, where we desire to send a message together with conveying information about Y^n . Another extension that may be of interest is to consider the case where we have distributed state information S_1 and S_2 which are correlated at two different action encoders. We are still interested in the output Y , but the additional dimension in this extension is in how the two distributed action encoders can coordinate to generate an output Y^n that is as compressible as possible.

VI. ACKNOWLEDGEMENT

This work is supported in part by the National Science Foundation (NSF) through Grant 0939370-CCF and in part by Air Force Office of Scientific Research (AFOSR) through Grant FA9550-10-1-0124. We thank Professor Abbas El Gamal for helpful discussion.

REFERENCES

- [1] C. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, Vol. 27, pp. 379423, 623656, 1948.
- [2] S. I. Gelfand and M. S. Pinsker, "Coding for Channel with Random Parameters," *Probl. Contr. and Inform. Theory*, Vol. 9, no. 1, pp. 1931, 1980.
- [3] C. Heegard and A. El Gamal, "On the Capacity of Computer Memory with Defects," *IEEE Trans. Inform. Theory*, Vol. IT-29, no. 5, pp. 731739, September 1983.
- [4] C. Shannon, "Channels with side information at the transmitter," *IBM J. Res. Develop.*, Vol. 2, pp. 289293, 1958.
- [5] O. Sumszyk, and Y. Steinberg, "Information embedding with reversible stegotext", in *Proceedings of the 2009 IEEE Symposium on Information Theory*, Seoul, Korea, Jun. 2009.
- [6] Y. H. Kim, A. Sutivong, and T.M. Cover, "State amplification," *IEEE Trans. Inform. Theory*, Vol. 54, no. 5, pp. 1850 -- 1859, May 2008.
- [7] A. Sutivong, and T. Cover, "Rate vs. Distortion Trade-off for Channels with State Information", in *Proceedings of the 2009 IEEE Symposium on Information Theory*, Lausanne, Switzerland, June 2002.
- [8] S. Sigurjonsson, and Y. H. Kim, "On multiple user channels with causal state information at the transmitters," in *Proceedings of IEEE International Symposium on Information Theory*, Adelaide, Australia, Sep. 2005.
- [9] Pulkit Grover, Aaron Wagner, and Anant Sahai, "Information Embedding meets Distributed Control", *IEEE Information Theory Workshop*, January 2010 in Cairo, Egypt.
- [10] R. F. Ahlswede, and J. Körner, "Source coding with side information and a converse for degraded broadcast channels," *IEEE Trans. Inf. Theory*, vol. 21, no. 6, pp. 629637, 1975.
- [11] A. El Gamal, and Y. H. Kim "Lectures on Network Information Theory," 2010, available online at ArXiv: <http://arxiv.org/abs/1001.3404>.
- [12] S. H. Lim, P. Minero, and Y.-H. Kim, "Lossy Communication of Correlated Sources over Multiple Access Channels," invited paper, Proceedings of the 48th Annual Allerton Conference on Communication, Control, and Computation, Monticello, Illinois, September 2010."
- [13] T.M. Cover and J.A. Thomas, "Elements of Information Theory," 2nd edition, Wiley Interscience Publication.
- [14] S. Jalali, A. Montanari and T. Weissman, "Lossy compression of discrete sources via Viterbi algorithm," *IEEE Trans. Inform. Theory*, Vol. 58, no. 4, April 2012.
- [15] S. Jalali and T. Weissman, "Block and sliding-block lossy compression via MCMC," accepted for publication in IEEE Transactions on Communications.

APPENDIX

A. Proof of Lemma 1

Fixing a v , the function $a = f(s, v)$ has only four possible forms: $a = s$, $a = 1 - s$, $a = 0$ and $a = 1$. Thus, we can divide \mathcal{V} into four groups:

$$\begin{aligned}\mathcal{V}_0 &= \{v : f(s, v) = s\} \\ \mathcal{V}_1 &= \{v : f(s, v) = 1 - s\} \\ \mathcal{V}_2 &= \{v : f(s, v) = 0\} \\ \mathcal{V}_3 &= \{v : f(s, v) = 1\}\end{aligned}\quad (13)$$

First, it is without loss of optimality to set $\mathcal{V}_3 = \emptyset$. That is because for each $v \in \mathcal{V}_3$, we can change the function to $f(s, v) = 0$. The rate $I(V; S) + H(Y|V)$ does not change and the cost EA only decreases.

Rewrite the objective function in the following way

$$\begin{aligned}I(V; S) + H(Y|V) &= H(Y|V) - H(S|V) + H(S) \\ &= H(S \oplus A \oplus Z|V) - H(S|V) + H(S) \\ &= \sum_{v \in \mathcal{V}_0} (H_2(p) - H(S|V = v))p(v) \\ &\quad + \sum_{v \in \mathcal{V}_1} (H_2(p) - H(S|V = v))p(v) \\ &\quad + \sum_{v \in \mathcal{V}_2} (H(S \oplus S_N|V = v) - H(S|V = v))p(v)\end{aligned}\quad (14)$$

where the last step is obtained by plugging in the actual form of $a = f(s, v)$ for each group of v .

Second, it is sufficient to have $|\mathcal{V}_0| = 1$ and $|\mathcal{V}_1| = 1$. To prove this, let $v_1, v_2 \in \mathcal{V}_0$. Note that $H(S|V = v)$ is a concave function in $p(s|V = v)$. Thus if we replace v_1, v_2 by a v_3 with $p(v_3) = p(v_1) + p(v_2)$ and

$$\begin{aligned}p(s|V = v_3) &= \frac{p(v_1)}{p(v_1) + p(v_2)}p(s|V = v_1) \\ &\quad + \frac{p(v_2)}{p(v_1) + p(v_2)}p(s|V = v_2),\end{aligned}$$

we preserve the distribution of S , the cost EA but we reduce the first term, i.e., $\sum_{v \in \mathcal{V}_0} (H_2(p) - H(S|V = v))p(v)$, in Eq. (14). Therefore, we can set $\mathcal{V}_0 = \{0\}$ and $\mathcal{V}_1 = \{1\}$.

Third note that for each $v \in \mathcal{V}_2$,

$$\begin{aligned}H(Y|V = v) - H(S|V = v) &= H(S \oplus A \oplus Z|V = v) - H(S|V = v) \\ &= H(S \oplus S_N|V = v) - H(S|V = v) \\ &\geq 0\end{aligned}\quad (15)$$

Last, if $\Pr(S = 0|V = 0) \neq \Pr(S = 1|V = 1)$, consider a new auxiliary random variable V' with the following distribution:

- $\mathcal{V}' = \{0, 1, 2\}$, $\Pr(V' = 0) = \Pr(V' = 1) = (\Pr(V = 0) + \Pr(V = 1))/2$
- The function $a = f(s, v')$ is of the form: $f(s, 0) = s$, $f(s, 1) = 1 - s$ and $f(s, 2) = 0$.
- $\Pr(S = 0|V' = 2) = 1/2$ and

$$\begin{aligned}\Pr(S = 1|V' = 0) &= \Pr(S = 0|V' = 1) \\ &= \frac{\Pr(S = 1|V = 0)\Pr(V = 0) + \Pr(S = 0|V = 1)\Pr(V = 1)}{\Pr(V = 0) + \Pr(V = 1)}.\end{aligned}$$

Comparing (S, V') with (S, V) , we can check that the cost EA and the distribution of S are preserved. Meanwhile, the objective function is reduced, which completes the proof.

B. Proof of Lemma 2

Similar to the proof of Lemma 1, we divide \mathcal{V} in to $\mathcal{V}_0, \mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3$. Using the same argument, we show that $\mathcal{V}_3 = \emptyset$.

Rewrite the objective function $H(Y|V)$ in the following way:

$$\begin{aligned}
 & H(Y|V) \tag{16} \\
 = & H(S \oplus A \oplus S_N|V) \\
 = & \sum_{v \in \mathcal{V}_0} H_2(p)p(v) \\
 & + \sum_{v \in \mathcal{V}_1} H_2(p)p(v) \\
 & + \sum_{v \in \mathcal{V}_2} (H(S \oplus S_N|V=v)p(v) \\
 = & H_2(p) \sum_{v \in \mathcal{V}_0 \cup \mathcal{V}_1} p(v) + \sum_{v \in \mathcal{V}_2} p(v),
 \end{aligned}$$

which implies that it is sufficient to consider the case $|\mathcal{V}_0| = 1$, $\mathcal{V}_1 = \emptyset$ and $|\mathcal{V}_2| = 1$. And this completes the proof.